

Evaluation of the Intel Xeon Phi and NVIDIA K80 as accelerators for two-dimensional panel codes[☆]

Lukas Einkemmer^{a,*}

^a*Department of Mathematics, University of Innsbruck, Austria*

Abstract

To predict the properties of fluid flow over a solid geometry is an important engineering problem. In many applications so-called panel methods (or boundary element methods) have become the standard approach to solve the corresponding partial differential equation. Since panel methods in two dimensions are computationally cheap, they are well suited as the inner solver in an optimization algorithm.

In this paper we evaluate the performance of the Intel Xeon Phi 7120 and the NVIDIA K80 to accelerate such an optimization algorithm. For that purpose, we have implemented an optimized version of the algorithm on the CPU and Xeon Phi (based on OpenMP, vectorization, and the Intel MKL library) and on the GPU (based on CUDA and the MAGMA library). We present timing results for all codes and discuss the similarities and differences between the three implementations. Overall, we observe a speedup of approximately 2.5 for adding a Intel Xeon Phi 7120 to a dual socket workstation and a speedup between 3 and 3.5 for adding a NVIDIA K80 to a dual socket workstation.

Keywords: Accelerators, GPU, Intel Xeon Phi, panel methods

1. Introduction

Numerical simulations are routinely used in applications to predict the properties of fluid flow over a solid geometry. Such applications range from the design and analysis of aircrafts to constructing more efficient wind turbines. In this context, a large number of different models and numerical methods have been developed in order to efficiently compute aerodynamic quantities such as lift and drag. It is generally accepted that the compressible Navier–Stokes system is able to represent the physics that is encountered in such systems faithfully. However, since solving the full Navier–Stokes equations is often computationally prohibitive, a hierarchy of models has been introduced that are able to provide good predictions for a selected range of problems.

In the following we will assume that the flow under consideration is irrotational and slow compared to the speed of sound (i.e., compressible effects can be neglected). This assumption reduces the Navier–Stokes equations to Laplace’s equation. One should note that a direct solution of Laplace’s equation would result in a body with zero lift. However, by imposing an additional constraint, the so-called Kutta condition, this simple model yields very accurate results in its regime of validity (even for lifting bodies such as airfoils, rotor blades, or fins). In addition, many phenomenological corrections have been developed that are able to extend the range of validity of this simplified model considerably.

In principle, any numerical method can be used to solve Laplace’s equation together with the Kutta condition. However, since we are usually interested in the fluid flow outside of a solid body, so-called panel methods (or boundary element methods) have become the standard approach. The advantage of such a method is that only the boundary has to be discretized. This implies that for a two-dimensional flow only a

[☆]This work is supported by the Austrian Science Fund (FWF) – project id: P25346.

*Corresponding author

Email address: lukas.einkemmer@uibk.ac.at (Lukas Einkemmer)

linear system in a single dimension has to be solved (although the corresponding matrix is no longer sparse). In addition, no error is made by introducing an artificial boundary faraway from the dynamics of interest. On modern computers a good implementation is able to compute, for example, the flow over an airfoil in less than a milliseconds (although this has not always been true in the past). Especially in the early days of computational fluid dynamics, performing such simulations was the only way to gain reasonable results. As a consequence, sophisticated software packages (such as Xfoil [5]) have been developed that are still used in current aerodynamics research (see, for example, [8, 9, 6]).

Panel methods are computationally cheap, and that fact makes them well suited as the inner solver in optimization algorithms. In this context the goal for the computer program is to find an ideal airfoil geometry given a target (or fitness) function. This is a nonlinear optimization problem as the geometry is the parameter under consideration. In recent years, genetic algorithms have enjoyed some success for such systems (see, for example, [1, 8, 7]). However, their application yields a new computational challenge as they require the computation of thousands or even hundred thousands of different airfoil configuration. This is a problem which lends itself very well to parallelization. As such it can potentially profit significantly from accelerators such as graphic processing units (GPUs) or the Intel Xeon Phi. In fact, a number of papers have been published that implement panel methods on GPUs (see, for example, the work conducted in [12, 2, 11, 3]). However, most of the literature focuses on the three dimensional case and the speedups reported seem to be somewhat optimistic and are almost never compared to a properly optimized central processing unit (CPU) code.

In this paper, we consider an implementation of a two-dimensional panel code that has been optimized for both the CPU and the Xeon Phi 7120 (using OpenMP and vectorization techniques) as well as for the NVIDIA K80 GPU (using CUDA). We use the Intel C++ compiler as well as the Intel MKL library which provides highly optimized LAPACK routines. In addition, we use the MAGMA library on the GPU (in order to further optimize the GPU code and exploit the second GPU present in the K80 expansion card).

The panel method used in this paper is introduced in section 2. We will discuss the hardware used and the corresponding performance characteristics in section 3. The timing results and details of the specific implementation under consideration are then presented in sections 4 (single GPU), 5 (Xeon Phi), and 6 (two GPU setup). Finally, we conclude in section 7.

2. Numerical algorithm

Panel methods are a type of boundary element methods and are widely used in fluid dynamics. In order to remedy the deficiency of Laplace's equation to describe the airflow over lifting bodies, they are supplemented by the empirically derived Kutta condition. This model, in many instances, gives a good description of lifting flow over solid bodies for low speed aerodynamics. In the following, we will limit ourselves to two-dimensional flows over wing cross sections (so-called airfoils).

The geometry of the problem is given by a sequence of points $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}^2$ that represent the discretization of an airfoil $\partial\Omega$. We assume that \mathbf{x}_0 is located at the trailing edge and that $\mathbf{x}_n = \mathbf{x}_0$ holds true. This setup is illustrated in Figure 1.




Figure 1: The discretized geometry of the NACA 2412 airfoil is shown (for the purpose of illustration a very coarse discretization with $n = 10$ is employed). The control points are shown in red and the exact geometry is outlined in gray.

The goal of the numerical method is to compute an approximation to the solution of Laplace's equation in $\mathbb{R}^2 \setminus \Omega$. The corresponding fundamental solution is given by

$$\phi(\mathbf{x}) = -\frac{1}{2\pi} \log |\mathbf{x}|.$$

Panel methods represent the solution as a superposition of fundamental solutions and the global flow that is imposed far away from the airfoil. That is, our solution $\varphi(\mathbf{x})$ will be written as

$$\varphi(\mathbf{x}) = \int_{\partial\Omega} \gamma(\mathbf{s}) \phi(\mathbf{x} - \mathbf{s}) d\mathbf{s} + \phi_{\mathbf{v}}(\mathbf{x}),$$

where $\gamma(\mathbf{s})$ is the coefficient in the superposition. The stream function of the global flow with velocity $\mathbf{v} = (v_1, v_2)$ is given by

$$\phi_{\mathbf{v}}(\mathbf{x}) = v_1 y - v_2 x = v_{\infty} y \cos \alpha - v_{\infty} x \sin \alpha.$$

The parameter α is called the angle of attack. Laplace's equation is subject to the boundary condition

$$\varphi|_{\partial\Omega} = C$$

which enforces that no fluid can move perpendicular to the wall. We discretize this ansatz by assuming that the vortex strength $\gamma(\mathbf{s})$ is constant on each panel. For a panel from \mathbf{x}_i to \mathbf{x}_{i+1} this yields

$$\begin{aligned} F_i(\mathbf{x}) &= \int_{\mathbf{x}_i}^{\mathbf{x}_{i+1}} \gamma_i \phi(\mathbf{x} - \mathbf{s}) d\mathbf{s} \\ &= \frac{\gamma_i}{2\pi} \frac{1}{|\mathbf{h}_i|} \left[\frac{1}{2} \langle \mathbf{x} - \mathbf{x}_i, \mathbf{h}_i \rangle \log |\mathbf{x} - \mathbf{x}_i|^2 - \frac{1}{2} \langle \mathbf{x} - \mathbf{x}_{i+1}, \mathbf{h}_i \rangle \log |\mathbf{x} - \mathbf{x}_{i+1}|^2 \right. \\ &\quad \left. - I \arctan 2(I, \langle \mathbf{x} - \mathbf{x}_i, \mathbf{h}_i \rangle) + I \arctan 2(I, \langle \mathbf{x} - \mathbf{x}_{i+1}, \mathbf{h}_i \rangle) - |\mathbf{h}_i|^2 \right] \end{aligned}$$

where $I = \langle \mathbf{h}_i^\perp, \mathbf{x} - \mathbf{x}_i \rangle$, $\mathbf{h}_i = \mathbf{x}_{i+1} - \mathbf{x}_i$, and \mathbf{h}_i^\perp is the outward pointing vector that is orthogonal to \mathbf{h}_i and satisfies $|\mathbf{h}_i^\perp| = |\mathbf{h}_i|$. We have used $\langle \cdot, \cdot \rangle$ to denote the dot product. The boundary condition is enforced at the control points (i.e., at $\mathbf{x}_{i+1/2} = (\mathbf{x}_{i+1} + \mathbf{x}_i)/2$). This yields an underdetermined system of linear equations

$$-\sum_{i=0}^{n-1} F_i(\mathbf{x}_{j+1/2}) + C = \sum_{i=0}^{n-1} A_{ji} \gamma_i + C = \phi_{\mathbf{v}}(\mathbf{x}_{j+1/2})$$

which we supplement by the Kutta condition

$$\gamma_0 = -\gamma_{n-1}.$$

In stating the Kutta condition we have assumed that the variables are ordered such that the trailing edge is located at $\mathbf{x}_0 = \mathbf{x}_n$. This, in total, gives us n equations for the n unknowns $\gamma_0, \dots, \gamma_{n-2}$ and C .

While the present numerical scheme yields good predictions for the lift coefficient, it gives completely wrong results for the drag coefficient. This is to be expected as drag is a viscous effect. However, a range of phenomenological corrections has been developed that, for attached flows, are able to predict the drag coefficient based on the inviscid solution. In our code we have implemented Thwaites' method (see, for example, [13, 10]) in order to perform a viscosity correction. Let us note, however, that more sophisticated schemes have been developed (see, for example, [5]). To validate our code we have run a genetic optimization algorithm based on a B-spline parametrization of the airfoils. In the genetic algorithm, tournament selection with one-point crossover is employed and the mutations are performed for a single B-spline coefficient at a time. The fitness function is proportional to the lift-to-drag ratio at zero angle of attack. A number of airfoils that were generated by our implementation are shown in Figure 2. The lift and drag coefficients stated are computed using Xfoil and confirm that our implementation generates successively better airfoils.

3. Hardware

The numerical implementation of the above algorithm requires two parts of significant computational effort. First, the system of linear equations has to be assembled which requires $\mathcal{O}(n^2)$ operations but involves the (expensive) evaluation of two logarithms and two arctan2 functions for each panel. Second,




Figure 2: Three airfoils for each generation of the genetic optimization algorithm are shown. The algorithm proceeds from the left to the right and each column represents a distinct generation. We show the best classes of airfoils (according to the lift-to-drag ratio) for a specific generation. The population size is equal to 1000.

the solution of the linear system of equations is usually done by an LU decomposition and thus involves $\frac{2}{3}n^3$ operations. In practice n is often between 100 and 300. In this regime both parts require substantial computational effort.

In the following we will consider the CPU, Xeon Phi, and GPU configuration listed in Table 1. These will be used for all the numerical simulations and all the benchmarks conducted in this paper. The corresponding (peak) performance characteristics with respect to single and double precision arithmetics and the theoretical attainable memory bandwidth are listed in Table 1. All of these components are part of a single dual socket workstation.

	TFlops/s		GB/s
	double	single	
1x E5-2630 v3	0.3	0.6	59
2x E5-2630 v3	0.6	1.2	59
1x Xeon Phi 7120	1.2	2.4	352
0.5x K80	1.5	4.4	240
1x K80	2.9	8.7	480

Table 1: Hardware characteristics of the dual socket workstation used in the numerical simulations. Peak arithmetic performance for single and double precision and the theoretically attainable memory bandwidth are listed.

Some fairly representative single and double precision timing results are collected in Table 2. These results point the clear picture that on the CPU assembling the matrix is between 2.5 and 3.5 times more expensive compared to solving the resulting linear systems. Thus, on the CPU the assembly actually dictates the performance of the algorithm to a large extend. This situation is reversed for both the Xeon Phi and the K80 GPU. For the Xeon Phi the assembly step is by approximately a factor of two faster compared to the two CPUs. Since assembly is an extremely compute bound problem and giving the similarities of the two architectures, this gain is expected based on the factor of two difference in the theoretical arithmetic performance. On the other hand, one half¹ of the K80 outperforms the same two CPUs by a factor of approximately 5.

It is striking that the performance of the linear solver is relatively poor on both the Xeon Phi as well as on the GPU. Note that in our application we are not interested in solving large linear systems (for which both of these libraries provide excellent performance) but in solving a large number of relatively small linear systems. Let us also note that the performance of the Intel MKL library on the CPU certainly suggests that there is some room for improvement for the Xeon Phi. In fact, some research has already been conducted in

¹Note that the NVIDIA K80 is a single expansion card that includes two identical GPUs each with its own separate memory. Thus, using one-half of the K80 means that we use one of the two GPUs.

single precision			
	Assembly	Solve	Total
E5-2630 v3	4.97	1.75	6.72
2x E5-2630 v3	2.76	1.07	3.83
Phi 7120	1.35	3.60	4.95
0.5x K80	0.46	3.70	4.16

double precision			
	Assembly	Solve	Total
E5-2630 v3	9.40	2.85	12.25
2x E5-2630 v3	5.19	2.05	7.24
Phi 7120	2.69	4.72	7.41
0.5x K80	0.79	4.42	5.21

Table 2: Time in seconds that is required to perform the assembly and linear solver step in our panel code. In the simulation 4000 candidate solutions (airfoil geometries) are optimized using a genetic algorithm with 10 generations. Each geometry is discretized using 200 points. For the linear solve we use the Intel MKL 2015 library (on the CPU and Intel Xeon Phi) and the MAGMA 1.7.0 linear algebra library on the GPU. All measured times are in units of seconds.

this direction (see, for example, [4, 14]). The same is presumably true for the GPU.

It thus seems that neither the CPU nor accelerators are ideally suited for the problem under consideration. However, since the accelerators are very efficient in the assembly step and the CPUs are very efficient in the linear solve step, the hope is that a hybrid algorithm that uses both platforms can succeed in obtaining a significant speedup compared to a CPU only implementation. The difficulty in this approach is that a large amount of data has to be transferred over the (relatively) slow PCIe bus. In the problem under consideration this means that all the assembled matrices have to be transferred from the accelerator to the CPU. Clearly, if such a scheme is to be successful some strategy has to be employed in order to mitigate this communication overhead. To present an efficient implementation and the corresponding benchmark results for both the Intel Xeon Phi and for the NVIDIA K80 is the purpose of the remainder of this paper.

4. GPU implementation

In this section we will consider an implementation where the assembly of the matrix is conducted exclusively on the GPU and the linear solves are performed exclusively on the CPU. This requires the transfer of a large number of matrices in each step from the GPU to the CPU. Timing results indicate that the run time of the assembly step (on the GPU) together with the required transfer of data (from the CPU to the GPU) is comparable or smaller than the time it takes to perform the linear solve (on the CPU). Thus, in order to hide the communication overhead, we interleave the assembly and transfer operation with the linear solves on the CPU. This approach is illustrated in Figure 3.




Figure 3: This figure shows a communication hiding pattern that interleaves the assembly (green; on the GPU) and copy (orange; data transfer from the GPU to the CPU) with the linear solve (blue; on the CPU). The red areas constitute the remaining overhead that decreases as we divide our problem into more and more slices.

Note that the overhead of this approach decreases as we increase the number of slices our problem is partitioned into. However, since the individual problems become smaller and smaller, overhead inherent in the different parts of the algorithm becomes more pronounced. Therefore, a compromise has to be made. In general, between 10 and 20 slices seems to yield near optimal performance in most circumstances.

single precision						
Hardware	slices	W	A	L	O	speedup
CPU	–	6.69	4.97	1.75	–	–
GPU+CPU	1	2.78	0.46	1.75	1.03	2.40
	5	2.12	0.46	1.87	0.25	3.16
	10	1.98	0.46	1.82	0.16	3.38
	20	1.97	0.46	1.86	0.10	3.40
2xCPU	–	3.80	2.76	1.07	–	–
GPU+2xCPU	1	2.14	0.47	1.09	1.06	1.78
	5	1.37	0.46	1.11	0.25	2.78
	10	1.26	0.46	1.11	0.16	3.01
	20	1.41	0.47	1.29	0.12	2.69

double precision						
Hardware	slices	W	A	L	O	speedup
CPU	–	12.21	9.40	2.85	–	–
GPU+CPU	1	4.82	0.77	2.90	1.92	2.53
	5	3.31	0.77	2.84	0.47	3.69
	10	3.13	0.77	2.84	0.29	3.91
	20	3.16	0.78	2.95	0.21	3.86
2xCPU	–	7.20	5.19	2.05	–	–
GPU+2xCPU	1	3.98	0.77	2.07	1.91	1.81
	5	2.63	0.77	2.15	0.48	2.73
	10	2.46	0.77	2.16	0.30	2.93
	20	2.50	0.78	2.28	0.22	2.88

Table 3: Timing results for the hybrid algorithm (GPU+CPU) illustrated in Figure 3. The wall time (W), the time required to assemble the system (A), the time required for the linear solves (L) and the overhead due to offloading to the GPU (O) are shown. The number of slices that yield the optimal run time are shown in bold. All measured times are in units of seconds.

The timing results are given in Table 3. We observe a speedup of 3 (single precision) and 2.9 (double precision) for adding a single K80 to the dual socket workstation. Although even a naive implementation (i.e., doing the assembly, the data transfer, and the linear solve in sequential order) results in some speedup, the communication hiding scheme employed contributes significantly to the performance of the implementation. In the case of a single socket workstation the observed speedup is approximately 3.4 (single precision) and 3.9 (double precision).

The overhead in this implementation can be partitioned into two parts:

- As we partition our problem into more and more slices the performance of the linear solver on the CPU decreases. This is a consequence of the overhead required for the asynchronous data transfer to the GPU as well as the overhead that is incurred in decreasing the batch size for the linear solver. In the numerical simulations conducted here this overhead is on the order of 10%.
- There is an inherent overhead in the interleave scheme (see the red area in Figure 3). This overhead decreases as we increase the number of slices.

Assuming instantaneous data transfer the optimal run time of our hybrid implementation is equal to the time for the linear solver. Our implementation is within 10 to 20% of that value.

5. Intel Xeon Phi implementation

In essence the implementation on the Xeon Phi is similar to the GPU implementation. However, there are two major differences. First, due to the 512 bit wide vector units, vectorization is extremely important to obtain good performance on the Xeon Phi. In order to enable the compiler to generate efficient code for the assembly step, we have added `__restrict` and `const` keywords to our code. We have used the vectorization report of the Intel C compiler to check that the compiler has indeed sufficient information to vectorize the time intensive portions of our algorithm. This has to be contrasted with the GPU implementation of the assembly step which is relatively straightforward (neither warp divergence nor coalesced memory access is a major concern in this application). Note, however, that the code for the Intel Xeon Phi is essentially identical to the optimized code for the CPU.

Second, since the assembly step takes significantly longer on the Xeon Phi compared to the GPU, it is no longer true that assembly (on the Xeon Phi) together with data transfer (from the Xeon Phi to the CPU) consumes less time than the linear solver (on the CPU). Thus, in order to obtain good performance we have to interleave all three operations as shown in Figure 4.




Figure 4: This figure shows a communication hiding pattern that interleaves the assembly (green; on the Xeon Phi), the copy (orange; data transfer from the Xeon Phi to the CPU), and the linear solve (blue; on the CPU). The red areas constitute the remaining overhead that decreases as we divide our problem into smaller and smaller slices.

The timing results for the Xeon Phi are given in Table 4. We observe a speedup of approximately 2.5 (for both single and double precision) for adding a single Xeon Phi 7120 to the dual socket workstation. On the other hand, for a single socket workstation the observed speedup is approximately 3 (single precision) and 3.6 (double precision).

Note that the performance of the GPU implementation considered in section 4 is superior by approximately 20% (for the dual socket case) and approximately 10% (for the single socket case) compared to the Xeon Phi implementation. We should also note that, as discussed before, the interleave scheme is out of necessity somewhat more complicated than the interleave scheme that is used for the GPU code (see Figure 3).

The performance difference between the Intel Xeon Phi and the GPU are mainly explained by the fact that the assembly step is more costly on the Xeon Phi. Therefore, it is not possible to hide the data transfer as well as on the GPU which negatively impacts the performance of the implementation.

6. Multiple GPU implementation

The GPU implementation in section 4 uses a single GPU to perform the assembly step of the optimization algorithm. However, as has been pointed out in the introduction, the NVIDIA K80 includes two identical GPUs within the same expansion card. Thus, so far we have only used one half of the computational potential within that package. Certainly, we can not expect a factor of two improvement when using this additional GPU as in the present implementation performance is mainly limited by the linear solve conducted on the CPU. However, the timing results given in Table 2 suggest that we could solve part of the problem (both assembly and linear solve) on the second GPU. In this situation, optimal load balancing dictates that about one quarter of the original problem is parceled out to the second GPU. Thus, in the best case we would expect an improvement in performance by approximately 30%.

There is one additional issue that deserves our attention. While the MAGMA linear algebra library includes routines that use the GPU memory as input and output, it is primarily designed to operate in an environment that includes CPUs as well as GPUs. Consequently, there is no way to execute a MAGMA

routine without CPU support and in an asynchronous fashion. To avoid oversubscription (which measurements show has a negative impact on performance) we use only 15 OpenMP threads for the linear solve and execute the MAGMA call in a separate pthread.

The timing results for this implementation are shown in Table 5. We observe a speedup of 3.1 (single precision) and 3.6 (double precision) for adding a K80 to the dual socket workstation. In the case of a single socket workstation the observed speedup is approximately 4.4 (single precision) and 5.1 (double precision). We remark that except for the dual socket single precision configuration we observe a speedup between 20 and 30% compared to the single GPU implementation.

7. Conclusion

We have compared the speedup that can be achieved when an Intel Xeon Phi 7120 or a NVIDIA K80 is added to a workstation with one or two Intel Xeon E5-2630 v3 CPUs. Optimization and parallelization for the CPU and Intel Xeon Phi code is done using the Intel C compiler (vectorization) and OpenMP. For the GPU we use an implementation that is based on CUDA. Since the linear solver is faster on the CPU and the assembly is faster on the Xeon Phi/GPU, the present algorithms profits from a hybrid implementation that uses both traditional CPUs as well as accelerators. The obtained results can be summarized as follows:

- Adding a K80 to the dual socket workstation results in a speedup of approximately 3.1 (single precision) and 3.6 (double precision).
- Adding a Xeon Phi 7120 to the dual socket workstation results in a speedup of approximately 2.4 (single precision) and 2.5 (double precision).
- Since the performance of the CPU only implementation is mostly dominated by the assembly step, the speedups for a single CPU are significantly larger. In this configuration we observe speedups of up to 5 for the GPU implementation and up to 3.5 for the Xeon Phi implementation.

From the results obtained we follow that for the problem under consideration the GPU implementation yields better performance compared to the Xeon Phi. What is not so clear cut is the development effort that is required for each platform. One advantage of the Xeon Phi is that once we had an optimized code for the assembly step on the CPU (using vectorization and OpenMP) we almost immediately obtained good performance on the Xeon Phi. On the other hand, the CUDA implementation of the assembly step is straightforward and due to the computational advantage of the GPU a less complicated communication hiding scheme proves sufficient. Thus, with respect to the development effort involved there is no clear winner.

References

References

- [1] V.D. Casás, R.J. Duro, and F. Lopez-Pena. Evolutionary design of wind turbine blades. *International Journal of Computing*, 4(3):49–55, 2014.
- [2] C.C Chabalko and B. Balachandran. Implementation and benchmarking of two-dimensional vortex interactions on a graphics processing unit. *Journal of Aerospace Information Systems*, 11(6):372–385, 2014.
- [3] N.B. deVelder. *Free Wake Potential Flow Vortex Wind Turbine Modeling: Advances in Parallel Processing and Integration of Ground Effects*. PhD thesis, University of Massachusetts-Amherst, 2014.
- [4] T. Dong, A. Haidar, P. Luszczek, J. Austin Harris, S. Tomov, and J. Dongarra. LU factorization of small matrices: accelerating batched DGETRF on the GPU. In *High Performance Computing and Communications, 2014 IEEE 6th Intl Symp on Cyberspace Safety and Security, 2014 IEEE 11th Intl Conf on Embedded Software and Syst (HPCC, CSS, ICESS)*, pages 157–160. IEEE, 2014.

- [5] M. Drela. XFOIL: An Analysis and Design System for Low Reynolds Number Airfoils. In T.J. Mueller, editor, *Low Reynolds Number Aerodynamics*, volume 54 of *Lecture Notes in Engineering*, pages 1–12. Springer Berlin Heidelberg, 1989.
- [6] O.S. Gabor, A. Koreanschi, and R.M. Botez. Low-speed aerodynamic characteristics improvement of ATR 42 airfoil using a morphing wing approach. In *IECON 2012–38th Annual Conference on IEEE Industrial Electronics Society*, pages 5451–5456. IEEE, 2012.
- [7] B.A. Gardner and M.S. Selig. Airfoil design using a genetic algorithm and an inverse method. *41st Aerospace Sciences Meeting and Exhibit*, 43, 2003.
- [8] B.R. Jones, W.A. Crossley, and A.S. Lyrantzis. Aerodynamic and aeroacoustic optimization of rotorcraft airfoils via a parallel genetic algorithm. *Journal of Aircraft*, 37(6):1088–1096, 2000.
- [9] T. Kipouros, T. Peachey, D. Abramson, and A. M. Savill. Enhancing and developing the practical optimisation capabilities and intelligence of automatic design software. In *8th AIAA Multi-Disciplinary Design Optimization Specialist Conference, Honolulu, Hawaii*, 2012.
- [10] J. Moran. *An introduction to theoretical and computational aerodynamics*. Dover, 2003.
- [11] G. Morgenthal, A. S. Corriols, and B. Bendig. A GPU-accelerated pseudo-3D vortex method for aerodynamic analysis. *Journal of Wind Engineering and Industrial Aerodynamics*, 125:69–80, 2014.
- [12] Y. Novikov, N. Sezer-Uzol, and O. Uzol. GPU based fast free-wake calculations for multiple horizontal axis wind turbine rotors. In *Journal of Physics: Conference Series*, volume 524, page 012100. IOP Publishing, 2014.
- [13] S. Tavoularis. Thwaites’ Method for Laminar Boundary Layers with Pressure Gradient (Lecture notes). http://by.genie.uottawa.ca/~mcg4345/AdditionalNotes/45_Thwaites.pdf.
- [14] A. Vladimirov. Fine-Tuning Vectorization and Memory Traffic on Intel Xeon Phi Coprocessors: LU Decomposition of Small Matrices. *Colfax Research*, 2015.

single precision						
Hardware	slices	W	A	L	O	speedup
CPU	–	6.69	4.97	1.75	–	–
Phi+CPU	1	3.70	0.97	1.72	1.98	1.80
	5	2.36	0.43	1.74	0.62	2.83
	10	2.25	0.27	1.81	0.44	2.97
	20	2.20	0.16	1.81	0.39	3.04
2xCPU	–	3.80	2.76	1.07	–	–
Phi+2xCPU	1	3.04	0.98	1.05	1.99	1.25
	5	1.77	0.42	1.11	0.67	2.14
	10	1.59	0.25	1.15	0.44	2.40
	20	1.65	0.18	1.22	0.43	2.31
GPU+CPU	20	1.97	0.46	1.86	0.10	3.40
GPU+2xCPU	10	1.26	0.46	1.11	0.16	3.01

double precision						
Hardware	slices	W	A	L	O	speedup
CPU	–	12.21	9.40	2.85	–	–
Phi+CPU	1	6.79	1.92	2.84	3.95	1.80
	5	3.90	0.81	2.73	1.17	3.13
	10	3.62	0.49	2.75	0.86	3.38
	20	3.43	0.28	2.77	0.66	3.56
2xCPU	–	7.20	5.19	2.05	–	–
Phi+2xCPU	1	5.96	1.92	2.01	3.95	1.21
	5	3.26	0.87	2.04	1.22	2.21
	10	2.97	0.49	2.10	0.87	2.42
	20	2.83	0.32	2.15	0.68	2.54
GPU+CPU	10	3.13	0.77	2.84	0.29	3.91
GPU+2xCPU	10	2.46	0.77	2.16	0.30	2.93

Table 4: Timing results for the hybrid algorithm (Xeon Phi+CPU) illustrated in Figure 4. The wall time (W), the time required to assemble the system (A), the time required for the linear solves (L) and the overhead due to offloading to the Phi (O) is shown. The number of slices that yield the optimal run time are highlighted in bold in the table. All measured times are in units of seconds.

single precision						
Hardware	distr	W	A	L	O	speedup
CPU	–	6.69	4.97	1.75	–	–
GPU+CPU	0.70	1.52	0.44	1.31	0.20	4.41
	0.75	1.58	0.45	1.44	0.15	4.22
	0.80	1.62	0.47	1.49	0.13	4.12
2xCPU	–	3.80	2.76	1.07	–	–
GPU+2xCPU	0.70	1.49	0.33	0.91	0.58	2.54
	0.75	1.29	0.35	0.95	0.34	2.94
	0.80	1.21	0.37	1.00	0.21	3.13
GPU+CPU	1.0	1.97	0.46	1.86	0.10	3.40
GPU+2xCPU	1.0	1.26	0.46	1.11	0.16	3.01

double precision						
Hardware	distr	W	A	L	O	speedup
CPU	–	12.21	9.40	2.85	–	–
GPU+CPU	0.70	2.44	0.55	2.20	0.23	5.01
	0.75	2.40	0.59	2.17	0.23	5.08
	0.80	2.66	0.62	2.42	0.23	4.60
2xCPU	–	7.20	5.19	2.05	–	–
GPU+2xCPU	0.70	2.01	0.55	1.70	0.31	3.57
	0.75	2.11	0.59	1.83	0.28	3.41
	0.80	2.26	0.62	2.00	0.26	3.18
GPU+CPU	1.0	3.13	0.77	2.84	0.29	3.91
GPU+2xCPU	1.0	2.46	0.77	2.16	0.30	2.93

Table 5: Timing results for the hybrid algorithm (GPUs+CPU) that uses both GPUs of the K80 hardware. The wall time (W), the time required to assemble the system (A), the time required for the linear solves (L) and the overhead due to offloading to the GPU (O) is shown. The number of slices that yield the optimal run time are highlighted in bold in the table. All measured times are in units of seconds.